

Huiwei Lu

Postdoc Appointee – Argonne National Laboratory
Phone: +1 (312) 647 7527 • Email: huiweilu@mcs.anl.gov
Homepage: www.mcs.anl.gov/~huiweilu

Education

Institute of Computing Technology <i>Doctor</i>	Computer Architecture 2007 – 2013
Beijing University of Chemical Technology <i>Bachelor</i>	Computer Science and Technology 2001 – 2005

Experiences

Argonne National Laboratory <i>Postdoc Appointee</i>	Parallel Runtime System 2013/8 – present
Argonne National Laboratory <i>Visiting Graduate Student</i>	High Performance Computing 2012/11 – 2013/4

Honors and Awards

- **Institute Chief Award** (Top 3%), Institute of Computing Technology, 2011
- Outstanding research assistant of the Computer Architecture Laboratory, Institute of Computing Technology, 2010
- Triple-A outstanding student of the Chinese Academy of Sciences, 2012, 2009, 2008
- Outstanding class leader of the Chinese Academy of Sciences, 2008

Professional Activities

- Program Committee: HPC China 2013
- Reviewer: The Journal of Supercomputing, Transaction of Cloud Computing, Cloudcom 2013
- External Reviewer: ISC 2014

Projects

Genome analysis at extreme scale 2012/11 – 2013/6
Design and implement a parallel k -mer matching library at scale.

- Define a programming model for overlap computation in genome assembly and mapping.
- Experiment on IBM Blue Gene/Q shows k -mer query throughput of 3.79×10^8 queries/s on 11.92 TB synthetic data with 65,536 cores.
- A *de novo* metagenomic assembler re-written using this library assembles 320 GB Human Microbiome Project sequence data in 10.9 minutes. $6.1 \times$ faster than the last version.

Trading computation for communication in large-scale graph traversal algorithm 2011/6 – 2013/6

Speedup $2.2\times$ and reduce communication time by 79.0% on a 512-node cluster.

- Compress the messages with bitmap compression, reducing the communication time by 52.4%.
- Sieve the messages with a novel distributed directory to make it more sparse, reducing the communication time by another 55.9%.
- Achieve a performance rate of 12.1 GTEPS (billion edge visits per second). Ranked 20th on the Graph 500 list, June 2012.

Tuning Linpack on Nebulae supercomputer

2012/1 – 2012/2

Achieve a performance of 620 teraflops with 2,000 nodes.

- Use *awk*, *grep* and *sed* to analyse system error logs, reducing the time of manually finding imbalance InfiniBand nodes among 2,000 nodes from hours to minutes.

Optimizing Graph 500 Benchmark

2011/4 – 2011/5

Speedup $1.9\times$ on a 32-node cluster.

- Use hybrid MPI/Pthreads programming to optimize the baseline application in Graph 500.
- Achieve a performance rate of 1.45 GTEPS. Ranked 18th on the Graph 500 list, June 2011.

Parallel Many-Core Processor Cluster Simulator

2010/2 – 2010/7

Achieve instruction-level simulation at a scale of 1,000 cores.

- Design and implement Network Interface Card simulation.
- Port MPICH to the simulator by creating a simulated device using the Channel Interface.

Parallel Many-Core Processor Simulator

2009/6 – 2010/1

Achieve an average speedup of $10.9\times$ (up to $13.6\times$), with accuracy lost of less than 0.1%.

- Use *parallel discrete event simulation* to parallelize a sequential many-core processor simulator which has 100,000 lines of code.

Publications

- Huiwei Lv, Guangming Tan, Mingyu Chen and Ninghui Sun. *Understanding Parallelism in Graph Traversal on Multi-core Clusters*. International Supercomputing Conference 2012 (ISC'12)
- Zehan Cui, Licheng Chen, Mingyu Chen, Yungang Bao, Yongbing Huang, Huiwei Lv. *Evaluation and Optimization of Breadth-First Search on NUMA Cluster*. IEEE Cluster 2012
- Huiwei Lv, Yuan Cheng, Lu Bai, Mingyu Chen, Dongrui Fan, and Ninghui Sun. *P-GAS: Parallelizing a Cycle-Accurate Event-Driven Many-Core Processor Simulator Using Parallel Discrete Event Simulation*. 24th ACM/IEEE/SCS Workshop on Principle of Advanced and Distributed Simulation (PADS 2010)
- Zheng Cao, Jianwei Xu, Mingyu Chen, Gui Zheng, Huiwei Lv, and Ninghui Sun. *HPPNetSim: A Parallel Simulation of Large-scale Interconnection Networks*. In Proceedings of the 2009 Spring Simulation Multiconference (SpringSim '09)
- Jianwei Xu, Mingyu Chen, Gui Zheng, Zheng Cao, Huiwei Lv, and Ninghui Sun. *SimK: A Large-Scale Parallel Simulation Engine*. Journal of Computer Science and Technology. November 06, 2009
- Huiwei Lv, Yuan Cheng, Lu Bai, Mingyu Chen, Dongrui Fan, Ninghui Sun. *Parallel Simulation of Many-core Processor and Many-core Clusters* (In Chinese). Journal of Computer Research and Development. 2013, Vol. 50

PhD Dissertation

Title: *On Optimizing Large-Scale Parallel Graph Algorithms*

Supervisors: Prof. Ninghui Sun

Technical Skills

MPI/Pthreads/OpenMP: Advanced

Erlang/CUDA/Hadoop: Beginner

C/Linux/Python/Shell: Advanced